arXiv:1511.07944v1 [stat.ML] 25 Nov 2015

Maximum Likelihood Estimation for Single Linkage Hierarchical Clustering

Dekang Zhu[a*]; Dan P. Guralnik[b]; Xuezhi Wang[c]; Xiang Li[a]; Bill Moran[c]

[a]*School of Electronic Science & Engineering, National University of Defense Technology, Changsha, China;* [b]*Electrical & Systems Engineering, University of Pennsylvania, Philadelphia, United States;* [c]*Electrical & Computer Engineering, RMIT University, Melbourne, Australia.*

We derive a statistical model for estimation of a dendrogram from single linkage hierarchical clustering (SLHC) that takes account of uncertainty through noise or corruption in the measurements of separation of data. Our focus is on just the estimation of the hierarchy of partitions afforded by the dendrogram, rather than the heights in the latter. The concept of estimating this "dendrogram structure" is introduced, and an approximate maximum likelihood estimator (MLE) for the dendrogram structure is described. These ideas are illustrated by a simple Monte Carlo simulation that, at least for small data sets, suggests the method outperforms SLHC in the presence of noise.

Keywords: Dendrogram; Ultra-metric; Minimum Spanning Tree; Markov Chain; Monte Carlo; Metropolis-Hastings algorithm.

1. Introduction

Distance-based clustering is the task of grouping objects by some measure of similarity, so that objects in the same group (or *cluster*) are more similar or closer (with respect to a prescribed notion of distance) than those in different clusters. Clustering is a common technique for statistical data analysis, widely used in data mining, machine learning, pattern recognition, image analysis, bioinformatics and cyber security.

Conventional ("flat", "hard") clustering methods accept a finite metric space (O, d) as input and return a partition of O as their output. Hierarchical clustering (HC) methods have a different philosophy: their output is an entire hierarchy of partitions, called a *dendrogram*, capable of exhibiting multi-scale structure in the data set [1, 2]. Rather than fixing the required number of clusters in advance, as is common for many flat clustering algorithms, it is more informative to furnish a hierarchy of clusters, providing an opportunity to choose a partition at a scale most natural for the context of the task at hand.

Many HC methods require linkage functions to provide a measure of dissimilarity between clusters (see [3, 4] for a fairly recent review). Some commonly used linkage functions are single linkage, complete linkage, average linkage, etc. The SLHC method, though suffering from the so called "chaining effect", remains popular for large scale applications [5] because of the low complexity of implementing it using minimum spanning trees (MST) [6]. This work relies on a particularly useful representation of dendrograms

*Corresponding author. Email: dekang.zhu@foxmail.com

using ultra-metrics, introduced by Jardine and Sibson [7]. Their point of view enabled redefinition of HC methods as maps from the collection of finite metric spaces to the collection of finite ultra-metric spaces [1, 8]. This enables a discussion of two essential properties — stability and convergence with respect to the Gromov-Hausdorff metric — that characterize SLHC within a broad class of HC methods [2].

Motivation: As described in [9], distance-based clustering methods, hierarchical as well as flat and overlapping, are deeply rooted in several mathematical disciplines, and are ubiquitous in bio-informatics applications. For example, in contemporary applications to the analysis of gene expression data [10–12], the raw data generated by microarrays is usually preprocessed to extract normalized expression values from which distance measures are computed, to be subsequently fed as input to a clustering algorithm. Depending on the kind of information sought, different variants of the conventional HC methods are applied, such as, for instance, hybrid HC [13] or improved Pearson correlation proximity-based HC [14].

More generally, HC methods play an important role wherever learning and analysis of data have to be performed in an unsupervised fashion. For example, clustering is a key underpinning technology in most algorithms for cyber-security. In this context, clustering arises in a large number of applications, including malware detection [15], identification of compromised domains in DNS traffic [16], classification of sources and methods of attacks [17], identification of hacker and cyber-criminal communities [18], detection of repackaged code in applications for mobile devices [19], and classification of files and documents [20]. There is an urgent need for more robust and reliable clustering algorithms.

Essentially all approaches to clustering, hierarchical or otherwise, accept the distances as "the truth". It is assumed uncorrupted by noise or artifacts. Particularly at the level of analogue data such as timing and device dependent parameters, but also even with some digital data, this is far from a correct model. For example, measures of dissimilarity between code samples are often engineered to reflect an opinion of the algorithm designer regarding the significance of specific features of executable code; it is more plausible to treat distance measures produced in this way as (quantifiably) uncertain measurements of the code sample rather than regard them as an objective truth.

Thus, practical necessities lead us to require that the output of clustering algorithms should account for uncertainty in the distance data and, to do that, a rigorous statistical approach is required. Obtaining dendrograms through statistical estimation (with an appropriate noise model for the data) will, in principle, result in improved HC methods to meet the needs of applications.

Conventional approaches to statistical estimation of partitions and hierarchies view the objects to be clustered as random samples of certain distributions over a prescribed geometry (e.g. Gaussian mixture model estimation using expectation-maximization in Euclidean spaces), and clusters can then easily be described in terms of their most likely origin. Thus, these are really *distribution-based* clustering methods — not distance-based ones. Our approach directly attributes uncertainty to the process of obtaining values for the pairwise distances rather than distort the data by mapping it into one's "favorite space". To the best of our knowledge, very little work has been done in this vein. Of note is [21], where similar ideas have been applied to the estimation of spanning trees in a communication network (see related work below).

Related work: In phylogenetic applications, the use of MLE and Bayesian methods for the estimation of evolutionary trees is a time-honored tradition spanning decades [22–29], and various clustering methods having been introduced for purposes of "phenetic clustering" [30, 31]. In a rough outline, one estimates a tree structure to describe a

population of samples from distributions of the form $p(\boldsymbol{x}|\tau)$, where τ is the evolutionary tree structure and the measurement $\boldsymbol{x}$ is a gene character (such as gene or nucleotide frequency); alternatively, one assigns (deterministically!) distance measures to reflect uncertainties in the quantities measured in the population. Estimation relying *only* on the noise model of the underlying dissimilarity measure would clearly constitute a much more general apparatus, compressing all the uncertainty about the data into the noise model, but otherwise treating all data sources with equal mathematical rigor.

A serious hurdle in the way of brute-force MLE estimation of dendrogram structure is the super-exponential growth of the number of such structures with the size of the data set. Naturally, this aspect of the estimation problem is more readily seen in applications related to cyber-space, where large data sets dominate the scene. The work of Castro *et al.* [21, 32] needs to be credited for having inspired an MCMC-based hypothesis-pruning procedure we have applied in this paper. However, we point out that both the clustering and the estimation problems that are the foci of their work are quite different from ours, and much more limited in scope. First, and most important, is that Castro *et al.* restrict attention to similarities with constant inter-cluster values, which effectively corresponds to postulating an ultra-metric setting *ab initio*. This is well-suited for the purposes of their application (network topology identification), but is unsatisfactory for the general case. Secondly, the network model of Castro *et al.* is not, strictly speaking a metric model, as they do not enforce the constraints coming from the triangle inequalities. The notorious complexity [33, 34] of this set of inequalities poses significant additional challenges to the problem of estimating structure from a measurement of a metric.

2. Preliminaries

Distance-Based Clustering. Given a set of objects O, a hard clustering method generates a partition of O — a collection of pairwise-disjoint subsets (*clusters*) of O whose union is O. For any $x \in O$ and any partition R of O, we denote the cluster of R containing x by R_x.

We focus on *distance-based clustering methods*, where a data set O undergoes initial processing to produce a symmetric, real-valued, non-negative function d on $O \times O$, whose values d_{xy} satisfy the triangle inequality, and serve to quantify the "degree of dissimilarity" between data entries. In applications, the user has some freedom to determine the values of the d_{xy} according to the requirements of the application in hand.

Hierarchical Clustering. Attempts have been made to anchor distance-based clustering in a firm axiomatic foundation, but results so far have been negative: [35] studies a seemingly intuitive and minimalistic axiomatic system for distance-based clustering which fails to support any clustering method; [36] works in a similar vein to demonstrate that reasonable axiomatic notions of distance between outputs of a flat clustering method are equally elusive. As later described in [1], the key obstruction to such axiomatic approaches lies with the requirement to produce a single partition as its output. To resolve this issue they proposed HC methods, producing *dendrograms*.

Recall that a partition R_1 is said to *refine* a partition R_2, if every cluster of R_1 is contained in a cluster of R_2; we denote this by $R_1 \succeq R_2$. A *hierarchy* is a collection $\mathcal{C}$ of partitions, where every $R_1, R_2 \in \mathcal{C}$ satisfy either $R_1 \succeq R_2$ or $R_2 \succeq R_1$. Intuitively, a dendrogram is a hierarchy with assigned heights, or *resolutions*; this is usually represented visually as a rooted tree — see Figure 1 (left).

Formally, following [2], we describe a dendrogram as a pair (O, β), where β is a map

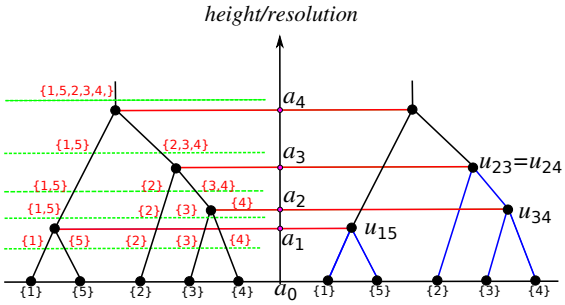

Figure 1. A rooted tree with labelled leaves as a dendrogram (left) and as an ultra-metric on $O = \{1, \ldots, 5\}$ (right)

of $[0, \infty)$ to the collection of partitions on O satisfying the following:

- There exists r_0 such that $\beta(r) = \{O\}$ for all $r \geqslant r_0$;
- If $r_1 \leqslant r_2$ then $\beta(r_1)$ refines $\beta(r_2)$;
- For all r, there exists $\epsilon > 0$ s.t. $\beta(t) = \beta(r)$ for $t \in [r, r + \epsilon]$.

Clusters of $\beta(r)$ are called *clusters at resolution* r.

Encoding Dendrograms. Ultra-metrics provide a convenient tool for encoding dendrograms [7]. Recall that a metric d on O is said to be an *ultra-metric*, if

$$d_{xy} \leqslant \max(d_{xz}, d_{zy}), \quad \forall\, x, y, z \in O. \tag{1}$$

The correspondence between dendrograms and ultra-metrics [2] is described as follows: any dendrogram β gives rise to an ultra-metric $u = u(\beta)$, as shown in Figure 1:

$$u(\beta)_{xy} := \inf \{r > 0 \mid \beta(r)_x = \beta(r)_y\}. \tag{2}$$

Conversely, the dendrogram β may be reconstructed from an ultra-metric u by setting

$$\beta(r)_x := \{y \in O \mid u_{xy} \leq r\}. \tag{3}$$

Single Linkage Hierarchical Clustering (SLHC). Single-linkage hierarchical clustering is defined, from the point of view of dendrograms, as follows. Given the metric space (O, d), for each $r \geq 0$, a dendrogram θ_d is constructed by setting $x, y \in O$ to lie in the same cluster of the partition $\theta_d(r)$ if and only if there exists a finite sequence of points $x_0, \ldots, x_m \in O$ with $x_0 = x$, $x_m = y$ and $d(x_{i-1}, x_i) \leq r$ for all $i \in \{0, \ldots, m\}$. Such a sequence is called an *r-chain from x to y in (O, d)*.

SLHC is often implemented by constructing an MST in (O, d). The partition $\theta_d(r)$ is obtained from *any* MST T of (O, d) by removing all edges of T of length$> r$; the clusters of the corresponding dendrogram $\theta_d(r)$ are the connected components of the resulting *forest*, and the corresponding ultra-metric distance u_{xy} $(x, y \in O)$ then equals the length of the longest edge of T (with respect to the distance d) separating x from y in T. Henceforth, we will write $u = \mathfrak{sl}(d)$ to denote the single linkage mapping of a metric d to the ultra-metric encoding of the corresponding dendrogram u.

It is a central result of [2], that SLHC is the unique hierarchical clustering method enjoying certain naturality properties, in stark contrast with the flat clustering situation. For a detailed discussion of the map $\mathfrak{sl}(\cdot)$, we refer the reader to [1, 2].

Notation. Note that metrics and ultra-metrics are conveniently written in matrix form, after fixing an order on O. Thus, writing $O := \{o_1, \ldots, o_n\}$, we will use $\boldsymbol{\theta} := [d_{ij}]$,

$d_{ij} := d(o_i, o_j)$ to denote the metric d in matrix form, and $\boldsymbol{u} = [u_{ij}]$ to denote the ultra-metric obtained from it by applying the map $\mathfrak{sl}(\cdot)$.

3. Statistical Model

It is useful to separate the metric information in a dendrogram (the grading by resolution) from the combinatorial information it conveys: a dendrogram/ultra-metric $\boldsymbol{u}$ is uniquely represented by a pair of parameters $(\tau, \boldsymbol{a})$, where:

- The parameter, τ, denotes the *structure of* u: the hierarchy defined by u (with the resolutions forgotten), ordered by refinement.
- The parameter, $\boldsymbol{a}$, is the *height vector of* u, whose coordinates, in order, indicate the minimum resolution at which each partition in the structure occurs — see Figure 1.

In what amounts to a choice of scale, we focus attention on the subset Θ of the space of all metrics θ satisfying $\mathfrak{sl}(\boldsymbol{\theta})_{ij} \leq 1$ for all $1 \leq i, j \leq n$. This is a compact convex set in $\mathbb{R}^{n(n-1)/2}$. Restricting attention to $\boldsymbol{\theta} \in \Theta$ is equivalent to placing a restriction on $\boldsymbol{a}$ to lie in the set Ω of all vectors $\boldsymbol{a}$ satisfying $0 \leqslant a_1 \leqslant a_2 \leqslant \cdots \leqslant a_{n-1} \leqslant 1$. Note that Θ coincides with the pre-image under $\mathfrak{sl}(\cdot)$ of the set of all ultra-metrics $\boldsymbol{u}$ with $\boldsymbol{a} \in \Omega$.

Remark 3.1 It must be observed that *degenerate structures*; that is, structures containing fewer than $n = |O|$ partitions (or, equivalently, corresponding to dendrograms that are not binary trees), occur in a set of metrics of Lebesgue measure zero, and therefore do not have any effect on statistical considerations regarding SLHC. Other clustering algorithms, such as hierarchical 2-means [37, 38], for example, do not have this property and, therefore, require more delicate analysis.

Our statistical model, introduced in [39], is as follows:

- The measurement $\boldsymbol{x}$ only depends on a metric $\boldsymbol{\theta} \in \Theta$ through a specific distribution $p(\boldsymbol{X}|\boldsymbol{\theta})$.
- The ultra-metric $\boldsymbol{u} = \mathfrak{sl}(\boldsymbol{\theta})$ is the parameter to be estimated from $\boldsymbol{x}$, with unknown $\boldsymbol{\theta}$ playing the rule of nuisance (latent) parameter.
- A reasonable assumption for this noise model is that the measurements of the different values of θ are sampled independently from the same parametrized distribution $p(X|\theta) = G_\theta(X)$, $\theta \in (0, +\infty)$:

$$p(\boldsymbol{X}|\boldsymbol{\theta}) = \prod_{ij} G_{\theta_{ij}}(X_{ij})\,; \tag{4}$$

Following the recommendations of [40], we pick integrated likelihood for our method of eliminating the nuisance parameter $\boldsymbol{\theta}$. Given a measurement $\boldsymbol{x}$ the likelihood function is:

$$\mathcal{L}(\boldsymbol{u}; \boldsymbol{x}) := p(\boldsymbol{x}|\boldsymbol{u}) = \int p(\boldsymbol{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}|\boldsymbol{u}) \mathbf{d}\boldsymbol{\theta} = \int \prod_{ij} G_{\theta_{ij}}(x_{ij}) p(\boldsymbol{\theta}|\boldsymbol{u}) \mathbf{d}\boldsymbol{\theta} \tag{5}$$

In the context of our problem the support of G_θ is restricted to $(0, +\infty)$.

The height vector $\boldsymbol{a}$ is implicit in the likelihood function, because of the complex integral, and so we focus on estimating the dendrogram structure τ, treating height vector $\boldsymbol{a}$ as a nuisance parameter for this task. It is reasonable to assume that the

structure and height functions are independent parameters. Replacing $\boldsymbol{u}$ by $(\tau, \boldsymbol{a})$, we obtain

$$p(\boldsymbol{x}|\tau) = \int p(\boldsymbol{x}|\boldsymbol{u})p(\boldsymbol{a})\mathrm{d}a = \int p(\boldsymbol{x}|\tau, \boldsymbol{a})p(\boldsymbol{a})\mathrm{d}a. \tag{6}$$

The MLE dendrogram structure is given by

$$\hat{\tau}(\boldsymbol{x}) = \arg\max_{\tau} p(\boldsymbol{x}|\tau). \tag{7}$$

3.1. *The likelihood of the metric given the data*

Since entries of $\boldsymbol{x}$ are measurements of distance, they are assumed to satisfy $x_{ij} > 0$ and, for the purposes of this paper, we assume each x_{ij} follows a log-normal distribution $\ln\mathcal{N}(\mu_{ij}, \sigma_{ij})$. Other noise distributions could be considered here. Therefore,

$$p(\boldsymbol{x}|\boldsymbol{\theta}) = \prod_{1\le i<j\le n} \frac{1}{\sigma_{ij}x_{ij}\sqrt{2\pi}} \exp\left(-\frac{(\ln x_{ij} - \mu_{ij})^2}{2\sigma_{ij}^2}\right), \tag{8}$$

where μ_{ij} and σ_{ij} are distribution parameters. We model the relationship between the measurement and the true metric $\boldsymbol{\theta}$ as follows: we require $\mathrm{E}[x_{ij}] = \theta_{ij}$ and $\mathrm{Var}[x_{ij}] = v$, which sets the distribution parameters accordingly to be

$$\mu_{ij} = \ln\theta_{ij} - \frac{1}{2}\ln(1 + \frac{v}{\theta_{ij}^2}), \quad \sigma_{ij} = \sqrt{\ln(1 + \frac{v}{\theta_{ij}^2})}. \tag{9}$$

3.2. *The likelihood of the dendrogram given the metric*

In the absence of assumptions on the latent parameters $\boldsymbol{\theta}$, $p(\boldsymbol{\theta}|\boldsymbol{u})$ is taken to be a uniform distribution on its support. This is justified by the fact that, considering the total weight/length of an MST as a natural energy functional on the space of metrics, it follows directly from (13) below that a maximum entropy distribution (subject to a total energy constraint) on this space restricts to a uniform distribution on the pre-image $\mathfrak{sl}^{-1}(\boldsymbol{u})$ of $\boldsymbol{u}$ under the single linkage map. Hence:

$$p(\boldsymbol{\theta}|\boldsymbol{u}) = \begin{cases} \frac{1}{\mathrm{Vol}(\mathfrak{sl}^{-1}(\boldsymbol{u}))} & \text{when } \boldsymbol{\theta} \in \mathfrak{sl}^{-1}(\boldsymbol{u}) \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

In more detail, following [39], denote

$$\mathrm{MST}(\boldsymbol{\theta}) = \{T \mid T \text{ is an MST of } \boldsymbol{\theta}\}. \tag{11}$$

For any spanning tree of K_O, the complete graph with vertex set O, write

$$C(T) := \{\boldsymbol{\theta} \mid \boldsymbol{\theta} \text{ is a metric on } O \text{ with } T \in \mathrm{MST}(\boldsymbol{\theta})\} \tag{12}$$

and observe that any two of the cones $C(T)$ intersect in a set of measure zero. In addition, the following identity is proved in [39]:

$$\mathfrak{sl}^{-1}(\boldsymbol{u}) = \bigcup_{T \in \mathrm{MST}(\boldsymbol{u})} C(T, \boldsymbol{u}) \tag{13}$$

where $C(T, \boldsymbol{u})$ is defined to be the set of all metrics $\boldsymbol{\theta}$ in $C(T)$ which coincide with $\boldsymbol{u}$ on the edges of T. In particular, any integration over $\mathfrak{sl}^{-1}(\boldsymbol{u})$ decomposes as a sum of integrals over the relevant $C(T, \boldsymbol{u})$. It is easy to verify that each $C(T, \boldsymbol{u})$ is a polytope given by the inequalities

$$C(T, \boldsymbol{u}) : \begin{cases} \theta_{ij} - \theta_{ik} - \theta_{kj} \leqslant 0 \\ -\theta_{ij} \leqslant -u_{ij} \\ \theta_{ij} \leqslant U_{ij} \end{cases} \tag{14}$$

where U_{ij} is defined to be the sum of the $\boldsymbol{u}$-lengths of all the edges separating o_i from o_j in the tree T for all i, j, k. Since membership in $C(T, \boldsymbol{u})$ is easily verified using the inequalities in (14), this enables Monte-Carlo evaluation of integrals over these domains.

3.3. *The Prior on the Space of Dendrograms*

We will assume that the parameter $\boldsymbol{a}$ is uniformly distributed in Ω. As a result,

$$p(\boldsymbol{a}) = \begin{cases} \dfrac{1}{\mathrm{Vol}(\Omega)} & \text{when } \boldsymbol{a} \in \Omega \\ 0 & \text{otherwise} \end{cases} \tag{15}$$

Thus, computing $p(\boldsymbol{x}|\tau)$ using Equation (6) requires the computation of integrals both over Ω and over the polytopes $C(T, \boldsymbol{u})$. Existing techniques for these computations, especially the latter, are computationally extremely complex [41–43]. In this paper we resort to numerical approximation techniques to be introduced in the next section.

4. Monte Carlo integration

Monte Carlo integration [44] is applied here to approximate the integrals.

4.1. *Equation* (5)

Setting up the Approximation. The inner integral $p(\boldsymbol{x}|\boldsymbol{u})$ splits, by (13), as a sum of integrals computed separately over each polytope $C(T, \boldsymbol{u})$. More precisely, we enumerate the MSTs associated with $\boldsymbol{u}$, $\mathrm{MST}(\boldsymbol{u}) = \{T_k\}_{k=1}^{K}$ where K may depend on $\boldsymbol{u}$, and set $C_k = C(T_k, \boldsymbol{u})$. Then $p(\boldsymbol{x}|\boldsymbol{u})$ becomes

$$p(\boldsymbol{x}|\boldsymbol{u}) = \int_{\mathfrak{sl}^{-1}(\boldsymbol{u})} p(\boldsymbol{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}|\boldsymbol{u}) \mathbf{d}\theta = \sum_{k=1}^{K} p(T_k|\boldsymbol{u}) \int_{C_k} p(\boldsymbol{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}|T_k) \mathbf{d}\theta. \tag{16}$$

Since $\boldsymbol{\theta}$ is uniformly distributed in $\mathfrak{sl}^{-1}(\boldsymbol{u})$, $p(\boldsymbol{\theta}|T_k)$ is uniform in C_k and $p(T_k|\boldsymbol{u})$ equals $\frac{\mathrm{Vol}(C_k)}{\mathrm{Vol}(\mathfrak{sl}^{-1}(\boldsymbol{u}))}$.

For the Monte-Carlo approximation we draw N samples $\{\boldsymbol{\theta}^{(l)}\}_{l=1}^{N}$ of metrics uniformly from $\mathfrak{sl}^{-1}(\boldsymbol{u})$, with N_k of these metrics drawn from each C_k. Then the Monte-Carlo approximation of $p(T_k|\boldsymbol{u})$ is N_k/N and, since $\sum_k N_k = N$, the Monte Carlo integration carried out for (16) yields the expression:

$$\phi(\boldsymbol{a}) := \sum_{k=1}^{K} \frac{N_k}{N} \cdot \frac{1}{N_k} \sum_{\boldsymbol{\theta}^{(l)} \in C_k} p\left(\boldsymbol{x}|\boldsymbol{\theta}^{(l)}\right) = \frac{1}{N} \sum_{l=1}^{N} p\left(\boldsymbol{x}|\boldsymbol{\theta}^{(l)}\right). \tag{17}$$

Drawing the Samples. According to Equation (14), lower and upper bounds on the value θ_{ij} of a metric $\boldsymbol{\theta} \in \mathfrak{sl}^{-1}(\boldsymbol{u})$ are, respectively,

$$\theta_{\min} = \min(\{u_{ij}\}), \quad \theta_{\max} = \max(\{U_{ij}\}). \tag{18}$$

Uniformly sampling metrics from the box defined by these bounds, and keeping only those lying in $\mathfrak{sl}^{-1}(\boldsymbol{u})$ as our samples for (17) is a straightforward approach, but fails to produce enough samples, because $\mathfrak{sl}^{-1}(\boldsymbol{u})$ has measure zero in the space of all nonnegative symmetric dissimilarity matrices in $\mathbb{R}^{n(n-1)/2}$. Instead, we draw metrics separately in each C_k, $k \in \{1, \ldots, K\}$ as described below and then combine these counts:

- **Step 1.** Set $\theta_{ij} = u_{ij}$ when edge $e_{ij} \in T_k$, then uniformly draw values between $\theta_{\min}$ and $\theta_{\max}$ for the other parameters;
- **Step 2.** Check the first constraint in (14) (triangle inequality) on each drawn vector, keep a certain number of these metrics;
- **Step 3.** Check the last two constraints in (14) for $T = T_k$, and keep the metrics satisfying them.

4.2. *Equation* (6)

Setting up the Approximation. The Monte Carlo approximation for the likelihood function $\mathcal{L}(\tau; \boldsymbol{x}) = p(\boldsymbol{x}|\tau)$ obtained from N_Ω samples $\{\boldsymbol{a}^{(l)}\}_{l=1}^{N_\Omega}$ drawn uniformly from Ω is:

$$\mathcal{L}(\tau; \boldsymbol{x}) \approx \int_{\boldsymbol{\Omega}} \phi(\boldsymbol{a}) p(\boldsymbol{a}) \mathbf{d}a \approx \frac{1}{N_\Omega} \sum_{l=1}^{N_\Omega} \phi(\boldsymbol{a}^{(l)}). \tag{19}$$

Drawing the Samples. Uniformly drawing a vector from Ω is not so straightforward a task. We observe that the uniform distribution on the standard simplex Δ is a Dirichlet distribution [45] with parameter vector $(1, 1, \ldots, 1)$. This can be mapped to Ω using the linear change of variables replacing the standard basis $(\boldsymbol{e}_i)_{k=1}^{n}$ of $\mathbb{R}^n$ by $\left\{\boldsymbol{w}_k := \sum_{i>k} \boldsymbol{e}_i\right\}_{k=1}^{n}$. Rewriting the vectors $\boldsymbol{a} \in \Omega$ in this basis yields:

$$\boldsymbol{a} = \sum_{i=1}^{n} \gamma_i w_i \,, \qquad \sum_{i=1}^{n} \gamma_i = 1, \gamma_i \in [0, 1]. \tag{20}$$

Equivalently, for any $\boldsymbol{a} = (a_k) \in \Omega$, we can write it as $a_k = \sum_{i=1}^{k} \gamma_i$ for $\boldsymbol{\gamma} = (\gamma_k) \in \Delta$. This transformation is volume preserving, so that uniform sampling of $\boldsymbol{a} \in \Omega$ is equivalent

to uniform sampling of $\boldsymbol{\gamma} \in \Delta$. This method produce samples of the required kind.

Algorithm 1 gives the pseudo-code for the numerical integration.

Algorithm 1 Pseudo-code for approximate computing of the likelihood (19)

function COMPUTE($\boldsymbol{x}$, τ);
  for $h = 1 : N_\Omega$ **do**;
    Draw a height vector $\boldsymbol{a}_h$ uniformly;
    $\boldsymbol{u}_h = (\tau, \boldsymbol{a}_h)$;
    Draw $\{\boldsymbol{\theta}^{(l)}\} \in \mathfrak{sl}^{-1}(\boldsymbol{u}_h)$ uniformly;
    Likelihood $prob(h) = \underset{l}{\text{mean}}\, p(\boldsymbol{x}|\boldsymbol{\theta}^{(l)})$;
  end for
  return $\log(\frac{\sum_h prob(h)}{N_\Omega})$.
end function

5. Reducing the Complexity

There are $n!(n-1)!/2^{n-1}$ elements in the set of combinatorial types of dendrograms on n particles [23]. Denote the space of all such types by Λ_n. The explosive growth of this set as a function of n makes brute-force maximization over it a prohibitive task even for reasonably small values of n. Nevertheless, our data (see Figures 2–5) indicates that structures of sufficiently low likelihood very rarely coincide with the target structure. The removal from consideration of such structures will result in little information loss for the outcome of MLE while significantly reducing computational cost. This suggests adaptation of a similar approach to that of [21] to produce an approximation of the MLE estimator: for a fixed measurement $\boldsymbol{x}$, we regard the likelihood $p(\boldsymbol{x}|\boldsymbol{\theta})$ as a distribution over Θ up to a normalizing factor denoted by η and draw a collection of metrics $\{\boldsymbol{\theta}_k\}$ from Θ using the Metropolis–Hastings (MH) algorithm [46] with target distribution $\eta p(\boldsymbol{x}|\boldsymbol{\theta})$. From the resulting collection of structures, we choose the subset $\Lambda_s \subset \Lambda_n$ of those structures appearing with highest frequencies, and run the computation from the previous section only for structures $\tau \in \Lambda_s$. (intuitively, as long as $\boldsymbol{x}$ is a measurement of reasonable quality, more metrics in the set $\{\boldsymbol{\theta}_k\}$ are close enough to the true metric $\boldsymbol{\theta}$ in order for $\boldsymbol{\theta}_k$ to support the same structure as $\boldsymbol{\theta}$ does).

For our implementation of the MH algorithm, we choose a proposal distribution

$$g(\boldsymbol{\theta}^{\text{old}} \to \boldsymbol{\theta}^{\text{new}}) = \prod_{ij=12}^{(n-1)n} \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{ -\frac{(\theta_{ij}^{\text{new}} - \theta_{ij}^{\text{old}})^2}{2\sigma^2} \right\} \tag{21}$$

with $\sigma = \sqrt{v}$, keeping in mind the possibility that the sample $\boldsymbol{\theta}^{\text{new}}$ might not be a metric, in which case the sample will be discarded (thus, $g(\boldsymbol{\theta}^{\text{old}} \to \boldsymbol{\theta}^{\text{new}})$ is, in fact, zero in the complement of Θ and is otherwise only proportional to the above expression).

As the proposal distribution is symmetric, for the acceptance probability we may use the Metropolis Choice:

$$A(\boldsymbol{\theta}^{\text{old}} \to \boldsymbol{\theta}^{\text{new}}) = \min\left(1, \frac{p(\boldsymbol{x}|\boldsymbol{\theta}^{\text{new}})}{p(\boldsymbol{x}|\boldsymbol{\theta}^{\text{old}})}\right) \tag{22}$$

We recall that samples are generated iteratively. At each iteration, a new state is drawn

from the proposal distribution for the current state. A real number q is drawn uniformly at random from $[0, 1]$, and the new state is accepted if $q \leq \frac{p(\boldsymbol{x}|\boldsymbol{\theta}^{\text{new}})}{p(\boldsymbol{x}|\boldsymbol{\theta}^{\text{old}})}$. Otherwise, the new state is rejected and the process remains in the same state. With additional burn-in and thinning, the iteration ends when a required number of metrics, N_θ, is obtained. Algorithm 2 summarizes the whole process.

Algorithm 2 Hypothesis pruning process for MLE estimation of dendrogram structure from a measurement $\boldsymbol{x}$ of a metric based on Metropolis-Hastings approximation of $\ell(\tau) \propto p(\boldsymbol{x}|\tau)$

function MH_SAMPLER($\boldsymbol{x}$)
  $\beta \leftarrow$ duration of burn-in period
  $\delta k \leftarrow$ thinning step
  $N_\theta \leftarrow$ number of metrics to be sampled from Θ
  $N_h \leftarrow$ number of hypotheses for output
  $\theta_0 \leftarrow$ arbitrary element of Θ
  for $k = 1$ **to** $\beta + \delta k \cdot N_\theta$ **do**
    $\theta_k \leftarrow$MH_TRANSITION(θ_{k-1})
  end for
  return N_h most frequently encountered structures
    from $\{\theta_k \mid k - \beta \geq 0\,,\ \delta_k$ divides $(k - \beta)\}$
end function

function MH_TRANSITION($\boldsymbol{\theta}^{\text{old}}$)
  repeat
    $\boldsymbol{\theta}^{\text{new}} \leftarrow$ a draw from (21)
  until $\boldsymbol{\theta}^{\text{new}} \in \Theta$
  $A \leftarrow \dfrac{p(\boldsymbol{x}|\boldsymbol{\theta}^{\text{new}})}{p(\boldsymbol{x}|\boldsymbol{\theta}^{\text{old}})}$
  $u \leftarrow$ draw from `Uniform(0, 1)`
  if $u \leq A$ **then**
    return $\boldsymbol{\theta}^{\text{new}}$
  else
    return $\boldsymbol{\theta}^{\text{old}}$
  end if
end function

6. Simulation

We demonstrate the effectiveness of the proposed hypothesis pruning process in the 5-particle case. A 5-point dendrogram may have any one of $5! \times 4!/2^4 = 180$ different dendrogram structures, which can be enumerated and indexed using the algorithm proposed in [29].

First, we randomly draw a structure from Λ_5 and a height vector from Ω to construct an objective ultra-metric (or, equivalently, a dendrogram). Then a random metric $\boldsymbol{\theta}^*$ is sampled from the pre-image (under $\mathfrak{sl}(\cdot)$) of this ultra-metric. This serves as the ground-truth metric used later to generate one measurement with a specified noise level. Finally, we implement MH as in Algorithm 2 with this measurement for its input to obtain a

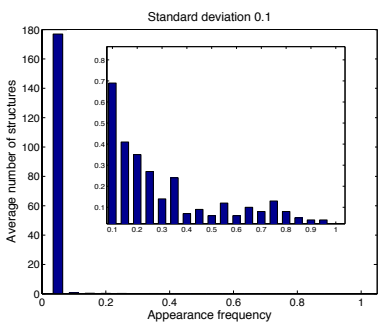

Figure 2. Average number of structures in each interval with std 0.1.

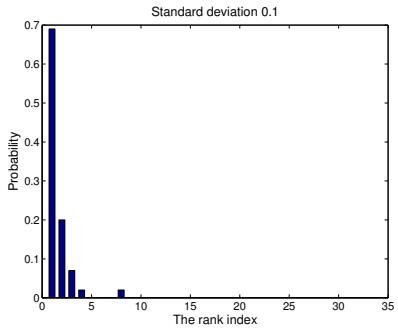

Figure 3. The probability distribution of the objective structure's rank with std 0.1.

sequence of metrics $\{\boldsymbol{\theta}_k\}$ from the distribution $\eta p(\boldsymbol{x}|\boldsymbol{\theta})$. For each measurement $\boldsymbol{x}$ of $\boldsymbol{\theta}^*$ 10000 steps of Algorithm 2 are computed, including $\beta = 1000$ steps of burn-in, and applying a thinning of $\delta_k = 3$ steps. The resulting output of $N_\theta = 3000$ observations is processed as follows.

(1) For each structure in Λ_5 its observed appearance frequency is calculated from among the $\{\boldsymbol{\theta}_k\}$ (the frequency is set to 0 for structures which did not arise).
(2) We subdivide the range $[0, 1]$ into 20 bins $[a_{i-1}, a_i]$ of equal lengths, and generate a vector $\boldsymbol{v}$ of 20 integers, indicating for each bin i the number of structures occurring with a frequency in $[a_{i-1}, a_i]$.
(3) This binning produces a ranking of the structures, in descending order, according to the frequency of occurrence.

Figure 2 shows the averaged histogram of the output vectors $\boldsymbol{v}$ generated from 1000 iid measurements $\boldsymbol{x}$ of $\boldsymbol{\theta}^*$ with standard deviation 0.1. Figure 3 shows the corresponding distribution of the rank of the true structure.

The inset in Figure 2 provides an enlarged plot of the bars excluding the leftmost one: observe that more than 170 of the 180 possible structures have insignificant appearance ratios (in the interval $[0, 0.05]$).

At the same time, Figure 3 shows that the rank of the true structure is almost completely distributed among the 10 highest ranked, with a nearly 70% chance of the true structure ranking first.

Figures 2 and 3 support our contention that most of the structures may be removed from consideration, and that with little chance of harm we may restrict attention to just a few of the highest ranking structures.

Figures 4 and 5 show our simulation results for noise with a standard deviation of 0.3. The majority of structures still appear with extremely low probabilities, though the rank of the true structure displays a more scattered distribution because of the increased noise.

To provide an idea of the overall performance of the proposed MLE estimator, Figure 6 compares the success rates of the MLE estimator with those of SLHC performed directly on the measurements. In this example, 16 height vectors were generated uniformly and, for each height vector, a dendrogram $\boldsymbol{u}$ with the indicated structure was created. For measurements, 1000 metrics were sampled from $\mathfrak{sl}^{-1}(\boldsymbol{u})$, and for each of them a measurement $\boldsymbol{x}$, drawn according to our data generation model with the prescribed standard deviation. For each $\boldsymbol{x}$, we first restricted attention to the most highly ranked 20 structures. Finally, the MLE estimator and SLHC were run on $\boldsymbol{x}$ and for each the frequency of

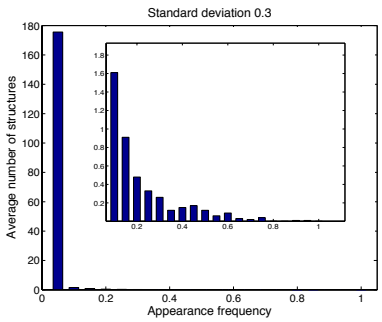

Figure 4. Average number of structures in each interval with std 0.3.

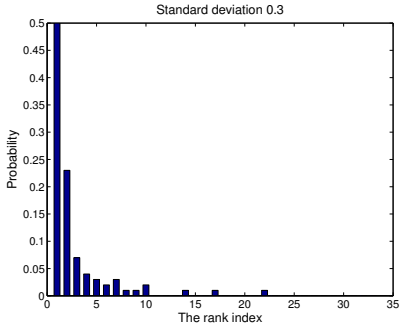

Figure 5. The probability distribution of the objective structure's rank with std 0.3.

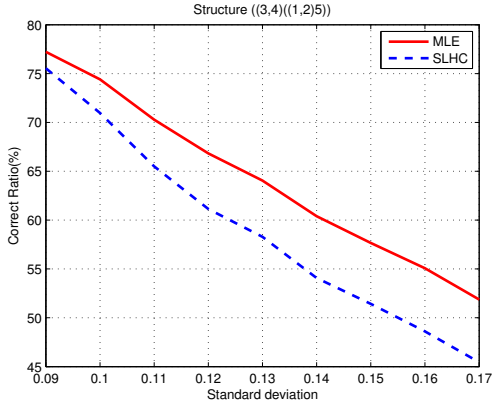

Figure 6. Comparison of MLE and SLHC

instances of successful identification of the initial $\boldsymbol{\tau}$ were recorded. These ratios were then averaged over the 16 heights. Figure 6 indicates that, on average, the proposed MLE estimator has a better error performance than SLHC, especially so for higher measurement noise variance.

7. Conclusion

This paper introduces a rigorous MLE approach to statistical estimation of SLHC under fairly general assumptions regarding the data generation process. Simulations with 5 particles demonstrate that the current approach used in all applications of SLHC — calculating SLHC directly from measured data — is significantly outperformed by the MLE estimation method.

A clear weakness of our current approach is its computational complexity which, as presented here, increases very rapidly with data size. This is largely because of the increased number of MSTs and the problem of sampling metrics in high dimensions, even though mitigated by our introduction of MCMC to cull the vast majority of structures. Further reducing the population of "MLE-eligible" MSTs is necessary, and suitable models are being investigated by us. It is likely that some variant on Kruskal moves, such as [47], will play a useful role here in providing a means to navigate spanning trees more effectively.

We also plan to consider a mixture of "top-down" and agglomerative methods of hierarchical approaches to further reduce the complexity of finding the "top split" in a computationally feasible way, and proceeding recursively from there, thereby reducing the search space in a step-by-step fashion.

Acknowledgements

This work was supported by the National Science Fund of China under Grant 61025006 and 61422114, and by the US Air Force Office of Science Research under a MURI FA9550-10-1-0567, and under FA9550-12-1-0418.

References

[1] Carlsson G, Mémoli F. Persistent clustering and a theorem of J. Kleinberg. arXiv preprint arXiv:08082241. 2008.

[2] Carlsson G, Mémoli F. Characterization, stability and convergence of hierarchical clustering methods. J Mach Learn Res. 2010 Aug;11:1425–1470; Available from: `http://dl.acm.org/citation.cfm?id=1756006.1859898`.

[3] Everitt BS, Landau S, Leese M, Stahl D. Cluster analysis (5th edition). 2011.

[4] Müllner D. Modern hierarchical agglomerative clustering algorithms. (preprint) http://arxivorg/abs/11092378. 2011;Available from: `http://arxiv.org/abs/1109.2378`.

[5] Jain AK, Murty MN, Flynn PJ. Data clustering: A review. ACM Comput Surv. 1999 Sep; 31(3):264–323; Available from: `http://doi.acm.org/10.1145/331499.331504`.

[6] Gower JC, Ross GJS. Minimum spanning trees and single linkage cluster analysis. Applied Statistics. 1969;54–64.

[7] N Jardine RS. Mathematical taxonomy. New York: John Wiley & Sons; 1971.

[8] Carlsson G, Mémoli F. Classifying clustering schemes. arXiv preprint arXiv:10115270. 2010.

[9] Sturmfels B. Can biology lead to new theorems. Annual report of the Clay Mathematics Institute. 2005;13–26.

[10] Shannon W, Culverhouse R, Duncan J. Analyzing microarray data using cluster analysis. Pharmacogenomics. 2003;4(1):41–52.

[11] Butte A. The use and analysis of microarray data. Nature reviews drug discovery. 2002; 1(12):951–960.

[12] Levenstien MA, Yang Y, Ott J. Statistical significance for hierarchical clustering in genetic association and microarray expression studies. BMC bioinformatics. 2003;4(1):62.

[13] Chipman H, Tibshirani R. Hybrid hierarchical clustering with applications to microarray data. Biostatistics. 2006;7(2):286–301.

[14] Booma P, Prabhakaran S, Dhanalakshmi R. An improved pearson's correlation proximity-based hierarchical clustering for mining biological association between genes. The Scientific World Journal. 2014.

[15] Wang J, Miller D, Kesidis G. Efficient mining of the multidimensional traffic cluster hierarchy for digesting, visualization, and anomaly identification. IEEE Journal on Selected Areas in Communications. 2006 Oct;24(10):1929–1941.

[16] Biggio B, Rieck K, Ariu D, Wressnegger C, Corona I, Giacinto G, Roli F. Poisoning behavioral malware clustering. In Proceedings of the 2014 Workshop on Artificial Intelligent and Security Workshop. ACM; 2014. p. 27–36.

[17] Kang J, Zhang Y, Ju J. Classifying DDoS attacks by hierarchical clustering based on similarity. In Machine Learning and Cybernetics 2006 International Conference; Aug; 2006. p. 2712–2717.

[18] Lu Y, Luo X, Polgar M, , Cao Y. Social network analysis of a criminal hacker community. Journal of Computer Information Systems. 2010;51(2):31–41.

[19] Hanna S, Huang L, Wu E, Li S, Chen C, Song D. Juxtapp: A scalable system for detecting code reuse among android applications. In Proceedings of the 9th International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment; Heraklion, Crete, Greece; DIMVA'12. Berlin, Heidelberg: Springer-Verlag; 2013. p. 62–81; Available from: http://dx.doi.org/10.1007/978-3-642-37300-8_4.

[20] Steinbach M, Karypis G, Kumar V. A comparison of document clustering techniques. In KDD Workshop on Text Mining; 2000.

[21] Castro RM, Member S, Coates MJ, Nowak RD. Likelihood based hierarchical clustering. IEEE Trans on Signal Processing. 2004;52:2308–2321.

[22] Cavalli-Sforza LL EA. Phylogenetic analysis: Models and estimation procedures. Am J Hum Genet. 1967 May;19:233–257.

[23] Edwards AW. Estimation of the branch points of a branching diffusion process. Journal of the Royal Statistical Society Series B (Methodological). 1970;155–174.

[24] Felsenstein J. Maximum likelihood and minimum-steps methods for estimating evolutionary trees from data on discrete characters. Systematic Biology. 1973;22(3):240–249; Available from: http://sysbio.oxfordjournals.org/content/22/3/240.abstract.

[25] Felsenstein J. Maximum-likelihood estimation of evolutionary trees from continuous characters. Am J Hum Genet. 1973;25:471–492.

[26] Felsenstein J. Evolutionary trees from dna sequences: A maximum likelihood approach. Journal of Molecular Evolution. 1981;17(6):368–376; Available from: http://dx.doi.org/10.1007/BF01734359.

[27] Yang Z. Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: Approximate methods. Journal of Molecular Evolution. 1994;39(3):306–314; Available from: http://dx.doi.org/10.1007/BF00160154.

[28] Rannala B, Yang Z. Probability distribution of molecular evolutionary trees: A new method of phylogenetic inference. Journal of Molecular Evolution. 1996;43(3):304–311; Available from: http://dx.doi.org/10.1007/BF02338839.

[29] Yang Z, Rannala B. Bayesian phylogenetic inference using DNA sequences: a Markov Chain Monte Carlo method. Molecular Biology and Evolution. 1997;14(7):717–724; Available from: http://mbe.oxfordjournals.org/content/14/7/717.abstract.

[30] Farris JS. Estimating phylogenetic trees from distance matrices. American Naturalist. 1972;645–668.

[31] Felsenstein J. Distance methods for inferring phylogenies: a justification. Evolution. 1984;16–24.

[32] Castro R, Nowak R. Likelihood based hierarchical clustering and network topology identification. Energy minimization methods in computer vision and pattern recognition. Springer Berlin Heidelberg; 2003.

[33] Deza A, Fukuda K, Pasechnik DV, Sato M. On the skeleton of the metric polytope. In Revised Papers from the Japanese Conference on Discrete and Computational Geometry; JCDCG '00. London, UK: Springer-Verlag; 2001. p. 125–136; Available from: http://dl.acm.org/citation.cfm?id=646320.686393.

[34] Deza MM, Laurent M. Geometry of cuts and metrics. Vol. 15 of Algorithms and Combinatorics. Springer, Heidelberg; 2010; first softcover printing of the 1997 original [MR1460488]; Available from: http://dx.doi.org/10.1007/978-3-642-04295-9.

[35] Kleinberg J. An impossibility theorem for clustering. Advances in neural information processing systems. 2003;463–470.

[36] Meilă M. Comparing clusterings: an axiomatic view. In Proceedings of the 22nd international conference on Machine learning. ACM; 2005. p. 577–584.

[37] Arslan O, Guralnik DP, Koditschek DE. Hierarchically clustered navigation of distinct euclidean particles. In 2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE; 2012. p. 946–953.

[38] Arslan O, Guralnik DP, Koditschek DE. Navigation of distinct euclidean particles via hierarchical clustering. In 2014 The Eleventh International Workshop on the Algorithmic Foundations of Robotics (WAFR2014); August; 2014.

[39] Zhu D, Guralnik DP, Wang X, Li X, Moran B. Statistical properties of the single linkage hierarchical clustering. arXiv preprint arXiv:1511.07715. 2015. Available from: http://arxiv.org/abs/1511.07715.

[40] Berger JO, Liseo B, Wolpert RL, et al. Integrated likelihood methods for eliminating nuisance parameters. Statistical Science. 1999;14(1):1–28.

[41] Barvinok AI. Computing the volume, counting integral points, and exponential sums. Discrete & Computational Geometry. 1993;10(1):123–141.

[42] Barvinok A, Hartigan J. Maximum entropy Gaussian approximations for the number of integer points and volumes of polytopes. Advances in Applied Mathematics. 2010;45(2):252–289.

[43] Lasserre JB, Zeron ES. A Laplace transform algorithm for the volume of a convex polytope. Journal of the ACM (JACM). 2001;48(6):1126–1140.

[44] MacKay DJ. Introduction to Monte Carlo methods. In Learning in graphical models. Springer; 1998. p. 175–204.

[45] Ng KW, Tian GL, Tang ML. Dirichlet and related distributions: Theory, methods and applications. Vol. 888. John Wiley & Sons; 2011.

[46] Chib S, Greenberg E. Understanding the Metropolis-Hastings algorithm. The American Statistician. 1995;49(4):327–335.

[47] Osipov V, Sanders P, Singler J. The filter-kruskal minimum spanning tree algorithm. In ALENEX. SIAM; 2009. p. 52–61; Available from: http://dblp.uni-trier.de/db/conf/alenex/alenex2009.html#OsipovSS09.